arXiv:1510.08178v2 [stat.ME] 6 Nov 2015

Clustering Via Finite Nonparametric ICA Mixture Models

Xiaotian Zhu*
Abbvie Inc.
and
David R. Hunter
Department of Statistics, Pennsylvania State University

November 28, 2015

Abstract

We propose an extension of non-parametric multivariate finite mixture models by dropping the standard conditional independence assumption and incorporating the independent component analysis (ICA) structure instead. We formulate an objective function in terms of penalized smoothed Kullback-Leibler distance and introduce the nonlinear smoothed majorization-minimization independent component analysis (NSMM-ICA) algorithm for optimizing this function and estimating the model parameters. We have implemented a practical version of this algorithm, which utilizes the FastICA algorithm, in the R package icamix. We illustrate this new methodology using several applications in unsupervised learning and image processing.

Keywords: independent component analysis, nonparametric estimation, penalized smoothed likelihood, unsupervised learning

*The authors gratefully acknowledge the support of Award No. DMS–1209007 from the National Science Foundation.

1 Introduction

Cluster analysis, or clustering, is one of several general approaches to the problem of unsupervised learning, that is, classification when no class labels are given (Hartigan 1975, Kaufman & Rousseeuw 2009). In practice, clustering is often based on heuristic ideas and intuitive measures that do not assume a probability model. As one example, hierarchical clustering (Hastie et al. 2009) builds a clustering tree, or dendrogram, using a linkage criterion that specifies dissimilarity between groups of observations. As another, k-means clustering (MacQueen et al. 1967) is centroid-based, where clusters are represented by a central vector and the algorithm aims at minimizing within-cluster sums of squares of distances.

By contrast, model-based clustering (Fraley & Raftery 1998, Banfield & Raftery 1993, Fraley & Raftery 2002) views data as coming from a mixture of probability distributions, each representing a cluster. Typically, the distributions are assumed to come from a parametric family such as normal, and group membership is learned from data by estimation algorithms which are often variations of the expectation-maximization method (Dempster et al. 1977). Since the early work of Wolfe (1963) and others, the literature on model-based clustering has expanded enormously. Indeed, if we view all work on finite mixture models as related to model-based clustering, then there are several book-length treatments of the subject, e.g., Titterington et al. (1985) and McLachlan & Peel (2004).

In recent years, with the rise of computational power, people have seen more studies focusing on semi-parametric or non-parametric model-based clustering; that is, methods that relax the standard model-based clustering assumption that the cluster densities come from a known parametric family. A semiparametric model-based clustering analysis for DNA microarray data can be found in Han & Davis (2006). Azzalini & Torelli (2007) proposes nonparametric density estimation using Delaunay triangulation for clustering via identification of subpopulations with regions with high density of the underlying probability distribution. Li et al. (2007) develops a clustering approach based on mode identification by applying new optimization techniques to a nonparametric density estimator. Vichi (2008) fits semiparametric clustering models to dissimilarity data. In Zhang et al. (2009), a semiparametric model is introduced to account for varying impacts of factors

over clusters by using cluster-level covariates. Mallapragada et al. (2010) proposes a nonparametric mixture model (NMM) for data clustering. Guglielmi et al. (2014) fits Bayesian semiparametric logit models to grouped data of in-hospital survival outcomes of patients hospitalized with ST-segment Elevation Myocardial Infarction diagnosis. Certain mixtures of linear regressions also fall under the category of semiparametric model-based clustering. For instance, Hunter & Young (2012) presents an algorithm for estimating parameters in a mixture-of-regressions model in which the errors are assumed to be independent and identically distributed but no other distributional assumption is made. Huang et al. (2013) proposes nonparametric finite mixture-of-regression models for analysis of U.S. housing price index (HPI) data. Vandekerkhove (2013) studies estimation of a semiparametric mixture-of-regressions model of two components when one component is known. Bajari et al. (2011) views a game abstractly as a semiparametric mixture distribution and studies the semiparametric efficiency bound of this model. Finally, Butucea & Vandekerkhove (2014) considers a semiparametric mixture of two distributions that are equal up to a shift parameter.

This article takes a different approach to nonparametric model-based clustering. As a starting point, we use recent advances (Zhu & Hunter 2015) on multivariate nonparametric finite mixture models under the conditional independence assumption. These models, which are described below in Section 2, have a much smaller literature and a more recent history that than of parametric mixture models: The basic idea was outlined by Hall & Zhou (2003), and various theoretical and algorithmic advances occurred gradually over the next decade, as documented in the survey article of Chauveau et al. (2015). The main novelty of the current article is to combine this existing work on nonparametric mixture models with the ideas independent components analysis (ICA) as elucidated, for example, by Hyvarinen et al. (2002). We explain the benefit of combining these ideas and justify the resulting combination theoretically. The result of our work is an algorithm we call the NSMM-ICA algorithm, which we implement in the **icamix** package for R (R Core Team 2015), available at `http://cran.r-project.org/web/packages/icamix/index.html`.

Previous work on model-based clustering using ICA has imposed parametric assumptions on the component density functions. Lee et al. (1999) and Lee et al. (2000) propose

parametric ICA mixture models with algorithms based on the infomax principle for various unsupervised classification problems. Shah et al. (2004) applies the ICA mixture model methodology to the problem of unsupervised classification of hyperspectral or multispectral imagery where image data are captured at multiple or a continuous range of frequencies across the electromagnetic spectrum. This is an important application of remote sensing and land cover classification. Palmer et al. (2008) derives an asymptotic Newton algorithm for Quasi-Maximum Likelihood estimation of the parametric ICA mixture model and presents its application to EEG segmentation.

We will present applications of our method to the analysis of several machine learning datasets, namely, the Fisher iris data, a dataset on Italian wines, a musical tone dataset due to Cohen (1984) (De Veaux 1989), and the learning of images in a dataset used by Lee et al. (1999) and several other authors.

2 The nonparametric ICA mixture model

Suppose that r-dimensional vectors $\mathbf{Y}_i = (Y_{i1}, Y_{i2}, ..., Y_{ir})^\top$, $1 \leq i \leq n$, are a simple random sample from a finite mixture of $m > 1$ components with positive mixing proportions $\lambda_1, \lambda_2, ..., \lambda_m$ that sum up to 1, and density functions $q_1, q_2, ..., q_m$ respectively. Here, m is known so the mixture density is

$$d(\mathbf{y}) = \sum_{j=1}^{m} \lambda_j q_j(\mathbf{y}) \tag{1}$$

for $\mathbf{y} \in R^r$. Equation (1) is often regarded as a semi-parametric model with $\lambda_1, \ldots, \lambda_m$ being the Euclidean parameters and $q_1, \ldots, q_m$ being the functional parameters.

An important special case of model (1) assumes the conditional independence of the coordinates of $\mathbf{y}$ given the component from which $\mathbf{y}$ is sampled. Under this assumption, each q_j equals the product of its marginal densities $q_{j1}, q_{j2}, ..., q_{jr}$ so that Equation (1) becomes

$$d(\mathbf{y}) = \sum_{j=1}^{m} \lambda_j \prod_{k=1}^{r} q_{jk}(y_k). \tag{2}$$

This assumption guarantees the generic identifiability of the parameters in the model for $r \geq 3$ under some weak assumptions (Allman et al. 2009). There exists a growing literature

on the theory and algorithmic treatment of model (1) under the conditional independence assumption. Chauveau et al. (2015) present a survey of this work, and Zhu & Hunter (2015) introduce some extensions that we will use in the current article.

Most previous work on model (2) assumes that we observe the random sample $\mathbf{Y}_1, \ldots, \mathbf{Y}_n$. However, in the current article we generalize this previous work by adding the assumption that the observed data are $\mathbf{X}_1, \ldots, \mathbf{X}_n$, where $\mathbf{X}_i = \mathsf{A}_j \mathbf{Y}$ for some invertible $r \times r$ matrix A_j, conditional on $\mathbf{Y}_i$ being generated from the jth component density q_j. In other words, we introduce additional parameters $\mathsf{A}_1, \ldots, \mathsf{A}_m$, one for each mixture component, consisting of the matrices that linearly transform the latent $\mathbf{Y}_i$ with independent coordinates into the observed $\mathbf{X_i}$. When there is no mixture structure, this assumption is exactly the independent component analysis (ICA) framework as described by Hyvarinen et al. (2002). NB: The word "component" in "ICA" is replaced by "coordinate" or "dimension" in the terminology of this article; here, "component" refers to one of the mixture densities.

To aid notation, let us define q_{A} for any nonnegative function q on R^r and invertible $r \times r$ matrix A as

$$q_{\mathsf{A}}(\mathbf{x}) = q(\mathsf{A}^{-1}\mathbf{x}) |\det \mathsf{A}|^{-1}, \tag{3}$$

which is the density function of a linearly transformed random variable having density q after left-multiplication by A.

Our ICA mixture model may thus be described formally as follows: We observe a random sample $\mathbf{X}_1, \ldots, \mathbf{X}_n$ from the mixture density

$$g(\mathbf{x}) = \sum_{j=1}^{m} \lambda_j f_j(\mathbf{x}), \tag{4}$$

where

$$f_j(\mathbf{x}) = (q_j)_{\mathsf{A}_j}(\mathbf{x}) \tag{5}$$

and

$$q_j(\mathbf{y}) = \prod_{k=1}^{r} q_{jk}(y_k). \tag{6}$$

For each observed $\mathbf{X}_i$, we shall define the corresponding latent variables

$$Z_{ij} = I\{ \mathbf{X}_i \text{ is drawn from the } j\text{th mixture component}\}$$

and $\mathbf{Y}_i = A_j^{-1}\mathbf{X}_i$ for the unique j such that $Z_{ij} = 1$. For estimation purposes, we write

$$(e_j)_{\mathsf{A}_j} = \lambda_j f_j, \tag{7}$$

so $e_j(\mathbf{x}) = \lambda_j q_j(\mathbf{x})$. To reduce ambiguities caused by over-parametrization, we further assume

$$E\{\mathbf{Y}_i \mathbf{Y}_i^\top\} = I_r, \tag{8}$$

where I_r is the $r \times r$ identity matrix. Assumptions (5), (6) and (8) are commonly used in the literature on ICA (Hyvarinen et al. 2002).

The q_{jk} do not follow any parametric form; any univariate density function in $L^1(R)$ is a valid candidate. Hence, we call the model described above a nonparametric ICA mixture model. This model is closely related to the ICA mixture model described in Lee et al. (2000), where each mixing component is parameterized by finitely many parameters. Here we relax this parametric assumption.

As a prelude to the next section on estimating the parameters of the nonparametric ICA mixture model, we define some operators that will aid notation. First, we define the smoothing, or convolution, operators S_h and S_h^*, Let $s_h(\cdot,\cdot) \in L^1(R \times R)$ be a nonnegative kernel function satisfying

$$\int s_h(v,z)\,\mathrm{d}z = \int s_h(v,z)\,\mathrm{d}v = 1 \tag{9}$$

for $v, z \in R$. Here, $h > 0$ is a user-specified tuning parameter often referred to as a bandwidth in smoothing contexts. For any $f \in L^1(R^r)$, let

$$(S_h f)(\mathbf{x}) = \int \tilde{s}_h(\mathbf{x},\mathbf{u}) f(\mathbf{u})\,\mathrm{d}\mathbf{u} \quad \text{and} \quad (S_h^* f)(\mathbf{x}) = \int \tilde{s}_h(\mathbf{u},\mathbf{x}) f(\mathbf{u})\,\mathrm{d}\mathbf{u}, \tag{10}$$

where

$$\tilde{s}_h(\mathbf{x},\mathbf{u}) = \prod_{k=1}^{r} s_h(x_k,u_k) \qquad \text{for } \mathbf{x},\mathbf{u} \in R^r. \tag{11}$$

Furthermore, let

$$(\mathcal{N}_h f)(\mathbf{x}) = \exp[(S_h^* \log f)(\mathbf{x})]. \tag{12}$$

Finally, we reproduce the projection-multiplication operator of Zhu & Hunter (2015), defined as

$$(Pf)(\mathbf{x}) = \frac{\left[\prod_{k=1}^{r} \int f(\mathbf{x})\,\mathrm{d}x_1\,\mathrm{d}x_2 \cdots \mathrm{d}x_{k-1}\,\mathrm{d}x_{k+1} \cdots \mathrm{d}x_r\right]}{\left[\int f\right]^{(r-1)}}. \tag{13}$$

Zhu & Hunter (2015) point out that when f is a density on R^r, the right side of (13) simplifies because the denominator is 1, and also that the P and S_h operators commute, i.e., $(P \circ S_h)f = (S_h \circ P)f$.

3 Parameter estimation

Now we consider the estimation of $\mathbf{e} = (e_1, e_2, ..., e_m)$ and $\mathsf{A} = (\mathsf{A}_1, \mathsf{A}_2, ..., \mathsf{A}_m)$. The idea is to use a penalized smoothed Kullback-Leibler divergence as the objective function for optimization, where we define the Kullback-Leibler divergence, KL, by

$$\mathrm{KL}\,(g_1, g_2) = \int \left[g_1 \log \frac{g_1}{g_2} + g_2 - g_1 \right] \tag{14}$$

for nonnegative functions g_1 and g_2 in $L^1(R^r)$. This is similar to the optimization function used in recent work of Levine et al. (2011) and Zhu & Hunter (2015), though the incorporation of the A_j into the usual conditional independence framework requires some delicacy. It is tempting to look at the Kullback-Leibler divergence between the target density and $\sum_j \mathcal{N}_h[(e_j)_{\mathsf{A}_j}](\mathbf{x})$. But applying the smoothing *after* the linear transformation ignores the different scale and direction of each coordinate of an observation from a mixture component. We believe a better approach is to compare the target density to $\sum_j [\mathcal{N}_h e_j]_{\mathsf{A}_j}(\mathbf{x})$ instead. Geometrically, the intuition behind this choice is that the spherical kernel should be applied to the data while they are approximately uncorrelated and standardized rather when they are on their original scale.

In the hypothetical case of a known target density $g(\mathbf{x})$, which we sometimes call the infinite sample size case, we therefore propose to minimize

$$\ell(\mathbf{e}, \mathsf{A}) = \int g(\mathbf{x}) \log \left[g(\mathbf{x}) / \sum_{j=1}^{m} [\mathcal{N}_h e_j]_{\mathsf{A}_j}(\mathbf{x}) \right] \mathrm{d}\mathbf{x} + \int \left[\sum_{j=1}^{m} (e_j)_{\mathsf{A}_j}(\mathbf{x}) \right] \mathrm{d}\mathbf{x}. \tag{15}$$

As in Zhu & Hunter (2015), we assume that for each $1 \le j \le m$, there exists $\theta_j > 0$ such that

$$e_j(\mathbf{x}) = \theta_j \prod_{k=1}^{r} e_{jk}(x_k), \tag{16}$$

where for each k, $1 \le k \le r$, $e_{jk} \in L^1(R)$ is positive. This overparameterization is employed for the sake of convenience and does not influence identifiability because we will never estimate θ_j separately.

Minimization of $\ell(\mathbf{e}, \mathsf{A})$ can be written equivalently as minimization of the penalized smoothed Kullback-Leibler divergence

$$KL\left(g, \sum_{j=1}^{m} (\mathcal{N}_h e_j)_{\mathsf{A}_j}\right) + \int \left[\sum_{j=1}^{m} (e_j)_{\mathsf{A}_j} - \sum_{j=1}^{m} [\mathcal{N}_h e_j]_{\mathsf{A}_j}\right](\mathbf{x})\, \mathrm{d}\mathbf{x}. \tag{17}$$

Note that the second term in (17) acts like a roughness penalty, and it is guaranteed to be positive since Jensen's inequality and Fubini's theorem give

$$\begin{aligned} \int [\mathcal{N}_h e_j]_{\mathsf{A}_j}(\mathbf{x})\, \mathrm{d}\mathbf{x} &\leq \iint |\mathsf{A}_j^{-1}| \tilde{s}_h(\mathbf{u}, \mathsf{A}_j \mathbf{x}) \exp \log e_j(\mathsf{A}_j \mathbf{x})\, \mathrm{d}\mathbf{u}\, \mathrm{d}\mathbf{x} \\ &= \int (e_j)_{\mathsf{A}_j}(\mathbf{x})\, \mathrm{d}\mathbf{x}. \end{aligned}$$

As Zhu & Hunter (2015) points out, a main advantage of using the $\mathbf{e}$ parameters instead of λ and $\mathbf{q}$ is that the latter parameterization requires the constraint that every q_j is a density function. With the $\mathbf{e}$ parameters, such a constraint is unnecessary because the minimizer of $\ell(\mathbf{e}, \mathsf{A})$ must be a density:

Theorem 3.1 *Any minimizer* $(\mathbf{e}^S, \mathsf{A}^S)$ *of (15) satisfies*

$$\int \sum_{j=1}^{m} e_j^S(\mathbf{x})\, \mathrm{d}\mathbf{x} = \int \sum_{j=1}^{m} (e_j^S)_{\mathsf{A}_j^S}(\mathbf{x})\, \mathrm{d}\mathbf{x} = 1. \tag{18}$$

Equation (18) demonstrates that for each $1 \leq j \leq m$, $\int e_j^S(\mathbf{x})\, \mathrm{d}\mathbf{x}$ can be interpreted as the corresponding mixing weight λ_j. The proof of Theorem 3.1 is essentially the same as the proof of Theorem 2.1 of Zhu & Hunter (2015).

4 The NSMM–ICA Algorithm

Here, we derive an iterative algorithm for solving the main problem of minimizing Equation (15). The algorithm is based on the MM framework, which stands for majorization-minimization (Hunter & Lange 2004) and which involves constructing and minimizing an alternative to the $\ell(\mathbf{e}, \mathsf{A})$ function at each iteration.

4.1 Majorizing the objective function

Given the current estimate $\mathbf{e}^{(0)}$ and $\mathsf{A}^{(0)}$, let us define

$$w_j^{(0)}(\mathbf{x}) = \frac{\left[\mathcal{N}_h e_j^{(0)}\right]_{\mathsf{A}_j^{(0)}}(\mathbf{x})}{\sum_{j'=1}^{m}\left[\mathcal{N}_h e_{j'}^{(0)}\right]_{\mathsf{A}_{j'}^{(0)}}(\mathbf{x})}. \tag{19}$$

Since $\sum_j w_j^{(0)}(\mathbf{x}) = 1$, Jensen's inequality gives

$$\begin{aligned}
&\ell(\mathbf{e},\mathsf{A}) - \ell(\mathbf{e}^{(0)},\mathsf{A}^{(0)}) \\
&= -\int g(\mathbf{x}) \log \sum_{j=1}^{m} w_j^{(0)}(\mathbf{x}) \frac{(\mathcal{N}_h e_j)_{\mathsf{A}_j}(\mathbf{x})}{(\mathcal{N}_h e_j^{(0)})_{\mathsf{A}_j^{(0)}}(\mathbf{x})}\,\mathrm{d}\mathbf{x} + \int\left(\sum_{j=1}^{m}(e_j)_{\mathsf{A}_j} - \sum_{j=1}^{m}(e_j^{(0)})_{\mathsf{A}_j^{(0)}}\right) \\
&\le -\int g(\mathbf{x}) \sum_{j=1}^{m} w_j^{(0)}(\mathbf{x}) \log \frac{(\mathcal{N}_h e_j)_{\mathsf{A}_j}(\mathbf{x})}{(\mathcal{N}_h e_j^{(0)})_{\mathsf{A}_j^{(0)}}(\mathbf{x})}\,\mathrm{d}\mathbf{x} + \int\left(\sum_{j=1}^{m}(e_j)_{\mathsf{A}_j} - \sum_{j=1}^{m}(e_j^{(0)})_{\mathsf{A}_j^{(0)}}\right).
\end{aligned}$$

Thus, if we let

$$b^{(0)}(\mathbf{e},\mathsf{A}) = -\int g(\mathbf{x}) \sum_{j=1}^{m} w_j^{(0)}(\mathbf{x}) \cdot \log\left(\mathcal{N}_h e_j\right)_{\mathsf{A}_j}(\mathbf{x})\,\mathrm{d}\mathbf{x} + \int\left(\sum_{j=1}^{m}(e_j)_{\mathsf{A}_j}\right),$$

then

$$\ell(\mathbf{e},\mathsf{A}) - \ell(\mathbf{e}^{(0)},\mathsf{A}^{(0)}) \le b^{(0)}(\mathbf{e},\mathsf{A}) - b^{(0)}(\mathbf{e}^{(0)},\mathsf{A}^{(0)}). \tag{20}$$

Therefore $b^{(0)}$ majorizes ℓ at $(\mathbf{e}^{(0)},\mathsf{A}^{(0)})$ up to an additive constant. We conclude that minimizing $b^{(0)}(\mathbf{e},\mathsf{A})$ will create an MM algorithm, as explained by Hunter & Lange (2004), and taking the next estimate in the iterative algorithm to be the minimizer will guarantee that the algorithm possesses a descent property.

4.2 Minimizing the majorizer

For each j, $1 \le j \le m$, we wish to minimize

$$b_j^{(0)}(e_j,\mathsf{A}_j) = -\int g(\mathbf{x}) w_j^{(0)}(\mathbf{x}) \cdot \log\left(\mathcal{N}_h e_j\right)_{\mathsf{A}_j}(\mathbf{x})\,\mathrm{d}\mathbf{x} + \int (e_j)_{\mathsf{A}_j}(\mathbf{x})\,\mathrm{d}\mathbf{x}. \tag{21}$$

Instead of finding a global minimizer for $b_j^{(0)}$, we first hold A_j fixed and minimize with respect to e_j, then plug in the resulting update to e_j and minimize with respect to A_j.

The resulting algorithm, which mimics the multiple "conditional maximization" steps of the ECM algorithm (Meng & Rubin 1993), does not actually minimize $b_j^{(0)}$, but it does ensure that the next iteration achieves a smaller value of $b_j^{(0)}$. This property is enough to guarantee the descent property, which states that the value of the objective function decreases at each iteration of the algorithm.

We find that Equation (21) has a closed-form minimizer as a function of e_j when A_j is held fixed.

Proposition 4.1 *The minimizer of Equation (21), with A_j held fixed, is*

$$\hat{e}_j(\mathbf{u}) = \frac{|\det \mathsf{A}_j|}{\left[\int g(\mathsf{A}_j\mathbf{x}) w_j^{(0)}(\mathsf{A}_j\mathbf{x})\ d\mathbf{x}\right]^{r-1}} \cdot \prod_{k=1}^{r} \int g(\mathsf{A}_j\mathbf{y}) w_j^{(0)}(\mathsf{A}_j\mathbf{y}) s_h(u_k, y_k)\, \mathrm{d}\mathbf{y}. \tag{22}$$

A proof of Proposition 4.1 is provided in Appendix A.

Equation (22) can be rewritten as

$$\hat{e}_j(\mathbf{u}) = \left[P \circ S_h(|\det \mathsf{A}_j| \cdot (g \cdot w_j^{(0)}) \circ \mathsf{A}_j)\right](\mathbf{u}) \tag{23}$$

using the P operator of Equation 13. In general, for any nonnegative function f on R^r,

$$S_h(f \circ \mathsf{A}_j) = (S_h)_{\mathsf{A}_j}(f) \circ \mathsf{A}_j, \tag{24}$$

where

$$(S_h)_{\mathsf{A}_j} f(\mathbf{x}) = \int |\det \mathsf{A}_j|^{-1} \tilde{s}_h(\mathsf{A}_j^{-1}\mathbf{x}, \mathsf{A}_j^{-1}\mathbf{u}) f(\mathbf{u})\, \mathrm{d}\mathbf{u}. \tag{25}$$

Thus, we may also write

$$(\hat{e}_j)_{\mathsf{A}_j}(\mathbf{u}) = \left[P_{\mathsf{A}_j} \circ (S_h)_{\mathsf{A}_j}(g \cdot w_j^{(0)})\right](\mathbf{u}), \tag{26}$$

where

$$P_{\mathsf{A}_j} f(\mathbf{u}) = [P(f_{\mathsf{A}_j^{-1}})]_{\mathsf{A}_j}(\mathbf{u}) = [P(f \circ \mathsf{A}_j)](\mathsf{A}_j^{-1}\mathbf{u}). \tag{27}$$

Now let us turn to the minimization with respect to A_j. We first define

$$\hat{q}_{jk}(u_k) = \frac{|\det \mathsf{A}_j|}{\int g(\mathsf{A}_j\mathbf{x}) w_j^{(0)}(\mathsf{A}_j\mathbf{x})\ \mathrm{d}\mathbf{x}} \int g(\mathsf{A}_j\mathbf{y}) w_j^{(0)}(\mathsf{A}_j\mathbf{y}) s_h(u_k, y_k)\, \mathrm{d}\mathbf{y}. \tag{28}$$

If we apply the change of variables $\mathbf{x} = \mathsf{A}_j\mathbf{y}$ to Equation (21) and then plug in $\hat{e}_j(\mathbf{u})$ into the resulting expression for $b_j^{(0)}(e_j, \mathsf{A}_j)$, we find that minimizing the result with respect to A_j is equivalent to minimizing

$$\log|\det \mathsf{A}_j| + \sum_{k=1}^{r} \int \hat{q}_{jk}(u) \log \hat{q}_{jk}(u)\, \mathrm{d}u \tag{29}$$

with respect to A_j, where $\hat{q}_{jk}$ depends on A_j through (28).

In Expression (29), $\hat{q}_{jk}$ is the kth margin of the kernel smoothed version of $(g \cdot w_j^{(0)})_{\mathsf{A}_j^{-1}} / \int g \cdot w_j^{(0)}$. In the discrete case where $dG(\mathbf{x})$ is the empirical distribution, $\hat{q}_{jk}$ is the kth margin of the kernel density estimate based on the linearly transformed (by A_j^{-1}) weighted observed data set, where the weight for the data point $\mathbf{x}_i$ is $w_j^{(0)}(\mathbf{x}_i)$. Let us denote this weighted data set by $\mathsf{D}_j^{(0)}$ and hence its linear transformation by $\mathsf{A}_i^{-1}\mathsf{D}_j^{(0)}$. By (28), the optimization mechanism at the current step views $\mathsf{A}_j^{-1}\mathsf{D}_j^{(0)}$ as a weighted sample generated from the unknown density function q_j, where $\mathsf{D}_j^{(0)}$ is a weighted sample from the jth mixing component and A_j^{-1} is the matrix that recovers the associated ICA transformations. Let us call A_j^{-1} a recovering matrix.

By Equation (8), we may treat $|\det A_j|$ as fixed given the weighted data $\mathsf{A}_j^{-1}\mathsf{D}_j^{(0)}$. The second term in (29) is an estimate of the sum of marginal entropies of q_j, which is equal, up to a term that does not involve A_j, to the mutual information of marginals of q_j. According to Hyvarinen et al. (2002), minimizing mutual information in this setting—that is, minimizing the mutual information of $\mathsf{A}_j^{-1}\mathbf{S}$ given a randomly chosen weighted sample from $\mathbf{S}$)—can be acheived by existing ICA algorithms such as the fastICA algorithm described in Section 4.3.

To summarize, the NSMM-ICA iterative algorithm will iterate as follows, where the parameters at the tth iteration will be denoted by $(\mathbf{e}^{(t)}, \mathsf{A}^{(t)})$:

Majorization Step: For $1 \le j \le m$, compute

$$w_j^{(t)}(\mathbf{x}) = \frac{(\mathcal{N}_h e_j^{(t)})_{\mathsf{A}_j^{(t)}}(\mathbf{x})}{\sum_{j=1}^{m} (\mathcal{N}_h e_j^{(t)})_{\mathsf{A}_j^{(t)}}(\mathbf{x})}. \tag{30}$$

ICA Step: Use ICA techniques to find $\mathsf{A}_j^{(t+1)}$ subject to (8) that minimizes

$$\sum_{k=1}^{r} \int \hat{q}_{jk}^{(t+1)}(u) \log \hat{q}_{jk}^{(t+1)}(u)\, \mathrm{d}u, \tag{31}$$

for $j = 1, \ldots, m$, where $\hat{q}_{jk}^{(t+1)}(u_k)$ is defined in Equation (28).

Minimization Step: Let

$$e_j^{(t+1)}(\mathbf{u}) = \hat{\lambda}_j^{(t+1)} \hat{q}_j^{(t+1)}(\mathbf{u}) = \hat{\lambda}_j^{(t+1)} \prod_{k=1}^{r} \hat{q}_{jk}^{(t+1)}(u_k), \tag{32}$$

where

$$\hat{\lambda}_j^{(t+1)} = \int (g \cdot w_j^{(t)}). \tag{33}$$

4.3 Practical Implementation of NSMM-ICA

Section 4.2 suggests alternating NSMM and ICA methods to form an iterative algorithm for the estimation of the nonparametric ICA mixture model. This section describes the practical considerations that went into the development of a package for R (R Core Team 2015), called **icamix**, that implements these ideas.

Empirical evidence suggests that NSMM and the npEM algorithm of Benaglia et al. (2009) tend to give very similar estimates (Levine et al. 2011). The reason is that usually $\mathcal{N}_h f$ is close to f itself. This suggests that the smoothed version of the algorithm can reasonably be replaced by the non-smoothed version because the former is more computationally burdensome than the latter. The decision to implement this non-smoothed version affects only Step 1 of the algorithm below. The result is an algorithm that fails to achieve the provable descent property of the smoothed version but which is much faster and which appears to result in nearly identical results for most test problems.

Among the many ICA techniques available in the literature, here we use the efficient and well-tested FastICA of Hyvarinen et al. (2002). At each iteration, FastICA will be applied to a weighted dataset, where the weight on observation i for component j is determined as the estimate, given the information available at the present iteration, of the probability that observation i falls into component j.

Assume we are given raw data as a matrix $\mathsf{X}^\top = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\}^\top$, where $\mathbf{x}_i = (x_{i1}, x_{i2}, x_{i3}, ..., x_{ir})^\top$ for $1 \leq i \leq n$. Our algorithm iterates through Steps 1 through 4 below until a convergence criterion is met.

Step 1. Estimate the jth component weight for the ith observation using the non-

smoothed densities:

$$p_{ij}^{(t)} = \frac{\lambda_j^{(t)} f_j^{(t)}(\mathbf{x}_i)}{\sum_{j\prime=1}^{m} \lambda_{j\prime}^{(t)} f_{j\prime}^{(t)}(\mathbf{x}_i)} = \frac{\lambda_j^{(t)} \left|\det \mathsf{A}_j^{(t)}\right|^{-1} \prod_{k=1}^{r} q_{jk}^{(t)} \left(\left[(\mathsf{A}_j^{(t)})^{-1} \mathbf{x}_i \right]_k \right)}{\sum_{j\prime=1}^{m} \lambda_{j\prime}^{(t)} \left|\det \mathsf{A}_{j\prime}^{(t)}\right|^{-1} \prod_{k=1}^{r} q_{j\prime,k}^{(t)} \left(\left[(\mathsf{A}_{j\prime}^{(t)})^{-1} \mathbf{x}_i \right]_k \right)}.$$

Step 2. Update the λ parameters:

$$\lambda_j^{(t+1)} = \frac{1}{n} \sum_{i=1}^{n} p_{ij}^{(t)}. \tag{34}$$

Step 3a. Centering FastICA step for component j:

$$\mathbf{x}_i \leftarrow \mathbf{x}_i - \frac{\sum_{i=1}^{n} \mathbf{x}_i p_{ij}^{(t)}}{\sum_{i=1}^{n} p_{ij}^{(t)}}. \tag{35}$$

Step 3b. Decorrelating FastICA step for component j: We first obtain the eigenvalue decomposition as

$$\frac{\sum_{i=1}^{n} \mathbf{x}_i \mathbf{x}_i^\top p_{ij}^{(t)}}{\sum_{i=1}^{n} p_{ij}^{(t)}} = EDE^\top, \tag{36}$$

then let $V = ED^{-1/2}E^\top$ and $\mathbf{z}_i = V\mathbf{x}_i$ for $i = 1, \ldots, n$. Therefore,

$$\frac{\sum_{i=1}^{n} \mathbf{z}_i \mathbf{z}_i^\top p_{ij}^{(t)}}{\sum_{i=1}^{n} p_{ij}^{(t)}} = VEDE^\top V^\top = I. \tag{37}$$

The transformed data $\mathbf{z}_i$ with weights $p_{ij}^{(t)}$, $1 \le i \le n$, are thus said to be "whitened"—i.e., their coordinates are uncorrelated and standardized—according to (37). Since $\mathbf{Z} = V\mathbf{X} = V\mathsf{A}_j\mathbf{S}$, we need to first estimate $(V\mathsf{A}_j)^{-1}$, of which the ith row is the same as $\mathbf{w}_i$ in (38) below, and multiply it by V on the right to get an update of A_j^{-1}.

Step 3c. Symmetric orthogonalization FastICA step for component j: The goal is to pick $\mathbf{w}_i \in R^r$ such that $\mathbf{w}_i^\top \mathbf{w}_i = 1$ for $i = 1, \ldots, r$. The first time we enter Step 3c, we may choose such $\mathbf{w}_i$ arbitrarily, for example setting $\mathbf{w}_i$ equal to the ith standard basis vector. At succeeding iterations, we begin an inner loop to modify the $\mathbf{w}_i$ from their previous values: We begin with

$$\mathbf{w}_i \leftarrow \frac{\sum_{i=1}^{n} \mathbf{z}_i g(\mathbf{w}_i^\top \mathbf{z}_i) p_{ij}^{(t)}}{\sum_{i=1}^{n} p_{ij}^{(t)}} - \mathbf{w}_i \frac{\sum_{i=1}^{n} g\prime(\mathbf{w}_i^\top \mathbf{z}_i) p_{ij}^{(t)}}{\sum_{i=1}^{n} p_{ij}^{(t)}}, \tag{38}$$

where g may be chosen to be either $g(y) = \tanh(\alpha_1 y)$ for some $1 \leq \alpha_1 \leq 2$ or $g(y) = y\exp(-y^2/2)$ (Hyvarinen et al. 2002). Denote $W = [\mathbf{w}_1, \mathbf{w}_2, ..., \mathbf{w}_r]^\top$ and then symmetrize and orthogonalize by

$$W \leftarrow (WW^\top)^{-1/2}W. \tag{39}$$

Iteratively update the $\mathbf{w}_i, i = 1, ..., r$ using Equation (38) and (39) until convergence is achieved. More precisely, we choose a tolerance τ and stop updating when

$$\max_{1\leq i\leq r}\left\{\left|\left(\mathbf{w}_i^{(previous)}\right)^\top \cdot \mathbf{w}_i^{(current)} - 1\right|\right\} \leq \tau. \tag{40}$$

Finally, set

$$\mathsf{A}_j^{(t+1)} = V^{-1}W^{-1}. \tag{41}$$

Step 4. Non-parametric density estimation step: Let

$$q_{jk}^{(t+1)}(\mathbf{u}) = \left(\sum_{i=1}^n p_{ij}^{(t+1)}\right)^{-1} \sum_{i=1}^n p_{ij}^{(t+1)} \frac{1}{h} K\left(\frac{\mathbf{u} - \left[(\mathsf{A}_j^{(t+1)})^{-1}\mathbf{x}_i\right]_k}{h}\right). \tag{42}$$

The R package **icamix** makes use of the **Rcpp** (Eddelbuettel & François 2011, Eddelbuettel 2013) and **RcppArmadillo** (Eddelbuettel & Sanderson 2014) packages for compiling and calling the core algorithms implementated in C++ code to speed up the calculations.

In the discrete algorithm we have developed, a single fixed bandwidth calculated from the data will not be sensible, especially because the scale is now changing according to the ICA framework. Thus we propose an iterative scheme for choosing the bandwidth similar to that of Benaglia et al. (2011), whereby

$$h_{jk}^{t+1} = 0.5 \cdot \mathrm{SD}_{jk}^{t+1} \cdot (n\lambda_j^{t+1})^{-0.2} = 0.5 \cdot (n\lambda_j^{t+1})^{-0.2}. \tag{43}$$

For simplicity, we replace $\min\{\mathrm{SD}_{j,l}^{t+1}, \mathrm{IQR}_{j,l}^{t+1}/1.349\}$ in the original Benaglia et al. (2011) formulation by $\mathrm{SD}_{j,l}^{t+1} = 1$. We also propose the ad hoc coefficient of 0.5 rather than Silverman's 0.9 used by Benaglia et al. (2011) in order to capture fine features of the density for better performance in the classification task. Our experience is that using 0.9 tends to oversmooth the estimated density. Simulation studies and applications we have run suggest that Equation (43) works well in practice.

When running the algorithm, we have determined that Step 4 dominates the computing time. By making use of Gaussian kernels and utilizing certain symmetric structure in evaluating some of Gaussians, we are able to lower the computing cost for the kernel density estimation step by about 50% with respect to the **mixtools** package. Further improvement should be possible via the Fast Gauss Transform (Raykar et al. 2005) and related techniques, though we have not implemented these improvements.

5 Applications

Here, we describe our experience applying the modified NSMM-ICA algorithm implemented in the **icamix** package to several datasets of varying size.

5.1 Iris Flower Classification

We first consider the well-known iris flower data set introduced by Fisher (1936). Edgar Anderson collected the data to quantify the morphologic variation of iris flowers of three related species: Iris setosa, Iris virginica and Iris versicolor. There are 50 samples from each of the three species. Four features were measured on each sample: the length and the width of the sepals and petals, in centimeters. One class is easily separated from the others, but the other two are relatively close to each other.

This data set has been popular as a benchmark test case for supervised learning, based on the linear discriminant model developed by Fisher (1936). Here, we view the data as unlabeled and apply our NSMM-ICA algorithm to it for unsupervised learning. Figure 1 shows a comparison of true species information and results from NSMM-ICA as well as k-means algorithms. We find that the NSMM-ICA achieves a quite stable classification error of 4.67%, while the results from k-means clustering are unstable and often inaccurate with an error rate greater than 30%.

5.2 Italian Wine Classification

The Italian wine data set is another popular data set used for comparing various classifiers (Forina et al. 1988, Aeberhard et al. 1992). It contains results of a chemical analysis of

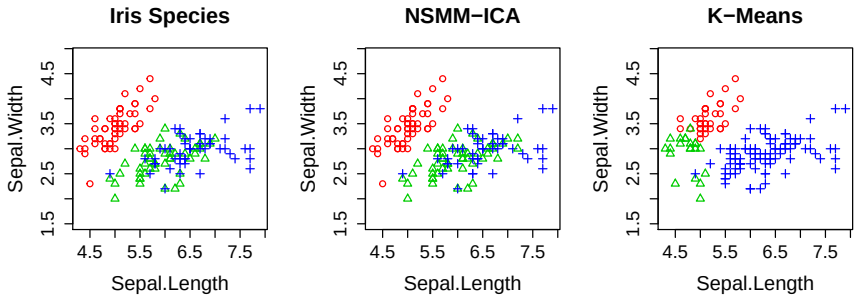

Figure 1: Iris Data: Comparison of true species information (far left) and two results from NSMM-ICA as well as k-means algorithms.

wines grown in Italy but derived from three different cultivars. A total of 178 observations are recorded, each with 13 continuous attributes such as color intensity, magnesium and malic acid. There are 59, 71 and 48 instances in the first (Barolo), second (Grignolino) and third (Barbera) wine classes, respectively.

Table 1: Wine Data classifications by PCA+NSMM-ICA algorithm

	Class 1	Class 2	Class 3
Barolo	59	0	0
Grignolino	6	61	4
Barbera	0	0	48

If we feed the unlabeled data directly to the NSMM-ICA algorithm, we obtain a classification error rate equal to 28.65%, prompting us to consider remedies. it seems that given the small number of observations, the relatively large number of attributes may be somewhat challenging as there are too many parameters to estimate and some of the attributes may consist of noise. So instead of using all 13 attributes, we first run principal component analysis (PCA) on the attributes and then select the 5 PCA scores that explain the largest proportion of variance in the attributes. Finally, we run the NSMM-ICA algorithm on the data set with the chosen PCA scores as attributes. In this way, the classification performance improves quite a lot, giving a classification error rate equal to 5.62%. Hence, in situations with relatively large numbers of coordinates, it might be worthwhile

to utilize a dimension reduction technique followed by the NSMM-ICA algorithm. Figure 2 shows a comparison of true species information and results from our unsupervised learning algorithms.

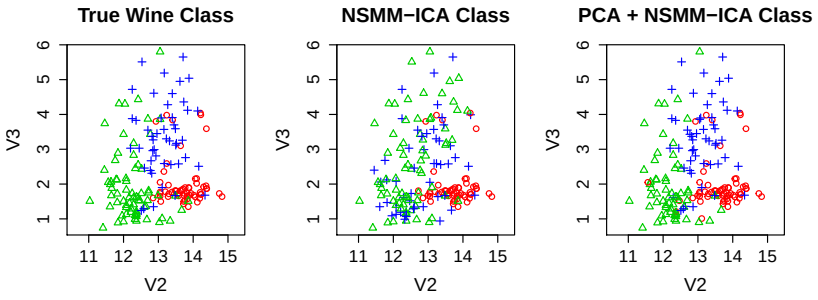

Figure 2: Wine Data: Comparison of true species Information (far left) and two results from our unsupervised learning algorithms.

5.3 Tone Data

The tone perception experiment and data were first introduced by Cohen (1984). These data have been analyzed by De Veaux (1989), Viele & Tong (2002), and Hunter & Young (2012) in the context of mixtures of regressions. In each trial of the experiment, a musician is presented with a fundamental tone plus a series of overtones determined by a stretching ratio. Then the musician is asked to tune an adjustable tone to one octave above the fundamental tone. Both the stretching ratio and the ratio of the adjusted tone to the fundamental are reported for each trial. There are five musicians involved in the experiment. However, the tone data set only contains 150 trials with the same musician. The problem of interest in conducting this experiment is to investigate the theory that the musician would either tune the tones to the nominal octave at a ratio of 2:1 to the fundamental tone (i.e., the interval memory hypothesis) or use the overtone to tune the tone to the stretching ratio (i.e., the partial matching hypothesis). The findings by Hunter & Young (2012) via modeling through a semi-parametric mixture of regressions conforms with the latter theory.

For this unsupervised learning task, we run both the npEM algorithm by Benaglia et al. (2009) and our NSMM-ICA algorithm on the tone data. Figure 3 shows that our NSMM-ICA algorithm does a good job of classification, very close to the mixture-of-regressions

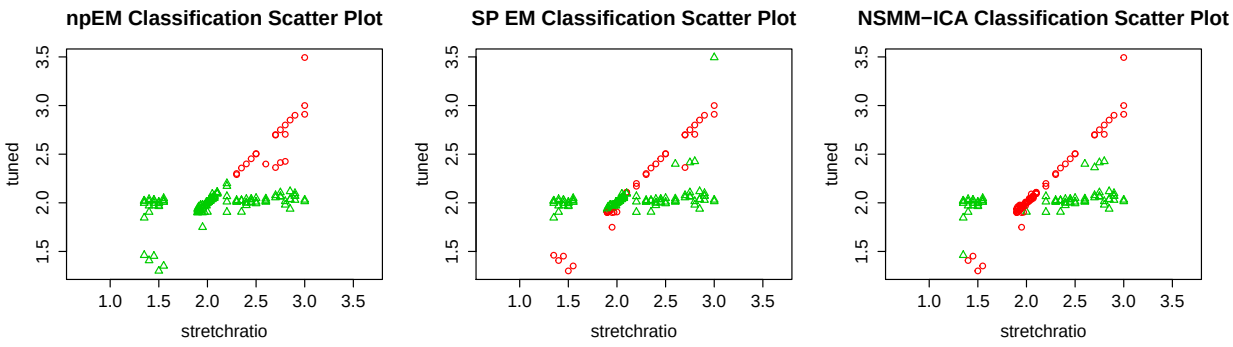

Figure 3: Comparison of three algorithms for fitting the tone dataset of Cohen (1984). Only the plot labeled SP EM explicitly assumes a mixture of regressions. The npEM algorithm does not utilize ICA and therefore misses the two lines entirely.

results obtained by Hunter & Young (2012), despite omitting any explicit assumption of regression structure. The reason why the results of the npEM algorithm shown in Figure 3 do not look nearly as good is because with regression lines that have nonzero slopes the mixture is far from being conditionally independent. Thus, the additional ICA step in our algorithm is essential.

For comparing our result with that of Hunter & Young (2012), we can use the estimated mixing weights obtained from NSMM-ICA to fit a weighted ordinary least squares model to obtain the regression coefficients. The results, summarized in Table 2, reflect the difference in estimated membership between SP EM and our NSMM-ICA primarily at the intersection of the two components, which is responsible for the noticeable difference in the estimated mixing weights.

5.4 Clustering images

Learning efficient codes for images obtained from different sources or contexts is an important problem in the area of image processing. The task involves extracting intrinsic structure in images by clustering and finding a complete set of efficient linear basis functions for each image source, which results in coefficient values being as statistically independent as possible. Techniques that utilize a parametric form of ICA mixture models have been proposed in Lee et al. (1999), Sejnowski (1999), and Lee et al. (2000). Here, we apply

Table 2: Comparison of mixtures of least squares fits for the tone dataset of Cohen (1984).

		SP EM	NSMM-ICA/Weighted LS
component 1	$\hat{\beta}_0$	1.77533	1.82215
	$\hat{\beta}_1$	0.11954	0.09076
component 2	$\hat{\beta}_0$	0.02121	-0.12111
	$\hat{\beta}_1$	0.97929	1.05584
$\hat{\lambda}_1$		0.67653	0.46779

the NSMM-ICA algorithm, which eliminates the parametric assumptions, to the task of unsupervised learning of image codes.

The two images shown in Figure 4 will be used as sources for the data set of the application. One is a painting image (2508× 1808 pixels) and the other is a newspaper image (2057×1365 pixels). The images are transformed to grey scale: Each pixel consists of a pixel intensity value ranging from 0 (black) to 1 (white). We select 5000 12×12 pixel

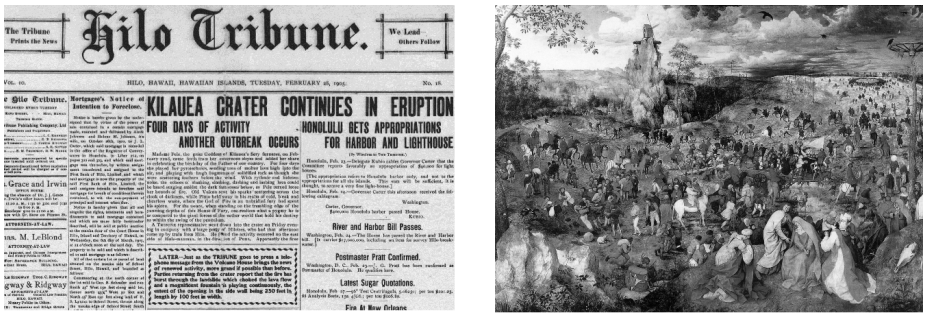

Figure 4: Two images as sources for data: Newspaper and painting.

patches randomly from each image. So the complete data set is of dimension 10,000×144. The NSMM-ICA algorithm converges after 19 iterations, which lasts a little less than 8 hours. Again each bandwidth is automatically learned by the iterative scheme we implemented. The result shows a very good recovery of the class-membership information, with a classification error rate of 1.2%.

The learned basis functions (i.e., a basis for the linear space of the pixel patches) show interesting patterns. Figure 5 shows the basis functions for each image. The ones for the

painting image appear smoother but more irregular, while the ones for the newspaper image look spottier and more regular.

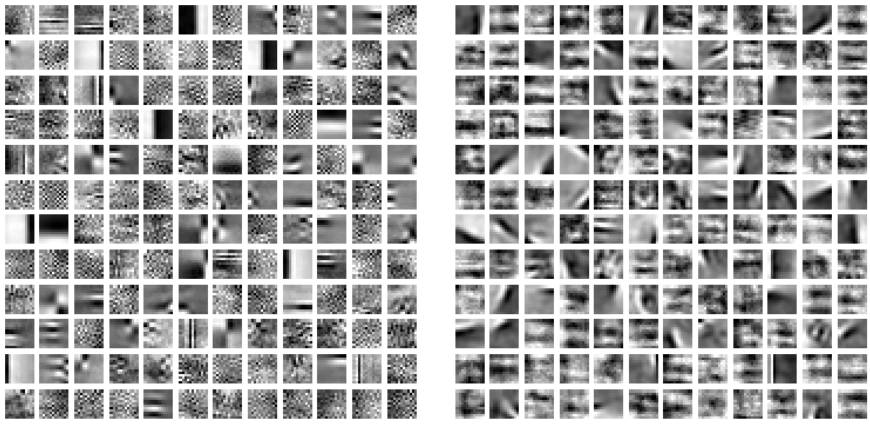

Figure 5: Learned basis functions for the newspaper (left) and painting (right). Each basis function, which is a 144-dimensional vector, is standardized to be within 0 and 1, then displayed as a 12×12 image patch.

6 Discussion

This article extends the work of Zhu & Hunter (2015) by loosening its conditional independence assumption, positing instead that each multivariate component can be linearly transformed to having independent coordinates. Clearly this gives a much more flexible model-based clustering framework than the conditional independence assumption alone. Among the particular favorable features of this extended model is that it allows for linear feature extraction, as illustrated by the tone data application of Section 5.3.

An optimization scheme similar in spirit to that under the original conditional independence model provides the theoretical justification for the NSMM-ICA algorithm, which combines the smoothed majorization-minimization method with independent component analysis (ICA) techniques for the estimation of the extended model. Yet additional theoretical questions remain. For instance, it is well-known that ICA cannot operate in the setting of multivariate normal data, since in that case the "whitened," or standardized, data is

already standard multivariate normal, so that there is no way to identify an ICA transformation. Thus, with the increased flexibility of our current framework come additional questions regarding conditions under which parameter identifiability holds. In addition, large-sample behavior of the NSMM estimator such as convergence rates is still not fully known (Chauveau et al. 2015), which means that of course this lack of knowledge extends to the current, more general, setting.

Due to computing efficiency considerations, practical implementation of the algorithm currently replaces the smoothed NSMM portion of the algorithm by the non-smoothed npEM of Benaglia et al. (2009). The **icamix** package for R (R Core Team 2015), available on CRAN, interweaves npEM with a weighted version of the FastICA algorithms (Hyvarinen et al. 2002). The package also implements an automated and adaptive scheme for bandwidth selection that is based on the work of Benaglia et al. (2011). Further computing efficiencies may be attainable through the use of the Fast Gauss Transform.

The proposed methodology has been tested via the **icamix** package on a number of applications in the area of unsupervised classification, and the results so far, as illustrated by the four examples of Section 5, have been positive. Given its flexibility and hence wide applicability, we believe that the novel approach to model-based clustering presented here has the potential to be a useful alternative to existing approaches based on parametric mixtures or mixtures that assume conditional independence.

A Proof of Proposition 4.1

The change of variables $\mathbf{x} = \mathsf{A}_j\mathbf{y}$ transforms Equation (21) into

$$\begin{aligned} b_j^{(0)}(e_j, \mathsf{A}_j) = &- \int g(\mathsf{A}_j\mathbf{y}) w_j^{(0)}(\mathsf{A}_j\mathbf{y}) \cdot \log\left\{(\mathcal{N}_h e_j)(\mathbf{y}) |\det \mathsf{A}_j|^{-1}\right\} |\det \mathsf{A}_j| \, \mathrm{d}\mathbf{y} \\ &+ \int e_j(\mathbf{y}) \, \mathrm{d}\mathbf{y}. \end{aligned}$$

Ignoring the term involving $\log |\det \mathsf{A}_j|^{-1}$ since it does not involve e_j, we find that minimizing $b_j^{(0)}(e_j, \mathsf{A}_j)$ as a function of e_j with A_j fixed is equivalent to minimizing

$$- \int g(\mathsf{A}_j\mathbf{y}) w_j^{(0)}(\mathsf{A}_j\mathbf{y}) \int \bar{s}_h(\mathbf{u}, \mathbf{y}) \log e_j(\mathbf{u}) |\det \mathsf{A}_j| \, \mathrm{d}\mathbf{u} \, \mathrm{d}\mathbf{y} + \int e_j(\mathbf{u}) \, \mathrm{d}\mathbf{u}, \tag{44}$$

which by Equations (9), (11), and (16) equals

$$-\sum_{k=1}^{r}\iint g(\mathsf{A}_j\mathbf{y})w_j^{(0)}(\mathsf{A}_j\mathbf{y})s_h(u_k,y_k)\log e_{jk}(u_k)|\det\mathsf{A}_j|\,\mathrm{d}u_k\,\mathrm{d}\mathbf{y}+\theta_j\int\prod_{k=1}^{r}e_{jk}(u_k)\,\mathrm{d}u_k \quad (45)$$

plus a term involving $\log\theta_j$ but none of the e_{jk}.

Picking a specific k and viewing Expression (45) as an integral with respect to $\mathrm{d}u_k$, we minimize it by minimizing the value of its integrand at each point. Differentiating with respect to $e_{jk}(u_k)$ and setting it equal to zero gives

$$-\frac{\int g(\mathsf{A}_j\mathbf{y})w_j^{(0)}(\mathsf{A}_j\mathbf{y})s_h(u_k,y_k)|\det\mathsf{A}_j|\,\mathrm{d}\mathbf{y}}{e_{jk}(u_k)}+\theta_j\left[\prod_{l\neq k}\int e_{jk}(u_l)\,\mathrm{d}u_l\right]=0,$$

yielding

$$\hat{e}_{jk}(u_k)\propto\int g(\mathsf{A}_j\mathbf{y})w_j^{(0)}(\mathsf{A}_j\mathbf{y})s_h(u_k,y_k)\,\mathrm{d}\mathbf{y},$$

which implies

$$\hat{e}_j(\mathbf{u})=\alpha_j\prod_{k=1}^{r}\int g(\mathsf{A}_j\mathbf{y})w_j^{(0)}(\mathsf{A}_j\mathbf{y})s_h(u_k,y_k)\,\mathrm{d}\mathbf{y} \quad (46)$$

for some constant α_j. To find α_j, we plug (46) into (44), differentiate with respect to α_j, and set the result equal to zero to obtain

$$\alpha_j=\frac{|\det\mathsf{A}_j|}{\left[\int g(\mathsf{A}_j\mathbf{x})w_j^{(0)}(\mathsf{A}_j\mathbf{x})\,\mathrm{d}\mathbf{x}\right]^{r-1}}, \quad (47)$$

which implies Equation (22).

References

Aeberhard, S., Coomans, D. & De Vel, O. (1992), 'Comparison of classifiers in high dimensional settings', *Dept. Math. Statist., James Cook Univ., North Queensland, Australia, Tech. Rep* (92-02).

Allman, E. S., Matias, C. & Rhodes, J. A. (2009), 'Identifiability of parameters in latent structure models with many observed variables', *The Annals of Statistics* **37**(6A), 3099–3132.

Azzalini, A. & Torelli, N. (2007), 'Clustering via nonparametric density estimation', *Statistics and Computing* **17**(1), 71–80.

Bajari, P., Hahn, J., Hong, H. & Ridder, G. (2011), 'A note on semiparametric estimation of finite mixtures of discrete choice models with application to game theoretic models', *International Economic Review* **52**(3), 807–824.

Banfield, J. D. & Raftery, A. E. (1993), 'Model-based gaussian and non-gaussian clustering', *Biometrics* pp. 803–821.

Benaglia, T., Chauveau, D. & Hunter, D. R. (2009), 'An EM-like algorithm for semi- and nonparametric estimation in multivariate mixtures', *Journal of Computational and Graphical Statistics* **18**(2), 505–526.

Benaglia, T., Chauveau, D., Hunter, D. R. et al. (2011), 'Bandwidth selection in an EM-like algorithm for nonparametric multivariate mixtures', *Nonparametric Statistics and Mixture Models: A Festschrift in Honor of Thomas P. Hettmansperger* pp. 15–27.

Butucea, C. & Vandekerkhove, P. (2014), 'Semiparametric mixtures of symmetric distributions', *Scandinavian Journal of Statistics* **41**(1), 227–239.

Chauveau, D., Hunter, D. R. & Levine, M. (2015), 'Semi-parametric estimation for conditional independence multivariate finite mixture models', *Statistics Surveys* **9**, 1–31.

Cohen, E. A. (1984), 'Some effects of inharmonic partials on interval perception', *Music Perception* pp. 323–349.

De Veaux, R. D. (1989), 'Mixtures of linear regressions', *Computational Statistics & Data Analysis* **8**(3), 227–245.

Dempster, A. P., Laird, N. M. & Rubin, D. B. (1977), 'Maximum likelihood from incomplete data via the EM algorithm', *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 1–38.

Eddelbuettel, D. (2013), *Seamless R and C++ Integration with Rcpp*, Springer, New York. ISBN 978-1-4614-6867-7.

Eddelbuettel, D. & François, R. (2011), 'Rcpp: Seamless R and C++ integration', *Journal of Statistical Software* **40**(8), 1–18.
URL: *http://www.jstatsoft.org/v40/i08/*

Eddelbuettel, D. & Sanderson, C. (2014), 'Rcpparmadillo: Accelerating r with high-performance c++ linear algebra', *Computational Statistics and Data Analysis* **71**, 1054–1063.
URL: *http://dx.doi.org/10.1016/j.csda.2013.02.005*

Fisher, R. A. (1936), 'The use of multiple measurements in taxonomic problems', *Annals of eugenics* **7**(2), 179–188.

Forina, M., Leardi, R., Armanino, C. & Lanteri, S. (1988), 'Parvus an extendible package for data exploration', *Classification and Correla* .

Fraley, C. & Raftery, A. E. (1998), 'How many clusters? which clustering method? answers via model-based cluster analysis', *The computer journal* **41**(8), 578–588.

Fraley, C. & Raftery, A. E. (2002), 'Model-based clustering, discriminant analysis, and density estimation', *Journal of the American Statistical Association* **97**(458), 611–631.

Guglielmi, A., Ieva, F., Paganoni, A. M., Ruggeri, F. & Soriano, J. (2014), 'Semiparametric bayesian models for clustering and classification in the presence of unbalanced in-hospital survival', *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **63**(1), 25–46.

Hall, P. & Zhou, X.-H. (2003), 'Nonparametric estimation of component distributions in a multivariate mixture', *Annals of Statistics* pp. 201–224.

Han, B. & Davis, L. S. (2006), Semi-parametric model-based clustering for dna microarray data, *in* 'Pattern Recognition, 2006. ICPR 2006. 18th International Conference on', Vol. 3, IEEE, pp. 324–327.

Hartigan, J. A. (1975), *Clustering algorithms*, Wiley.

Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J. & Tibshirani, R. (2009), *The elements of statistical learning*, Vol. 2, Springer.

Huang, M., Li, R. & Wang, S. (2013), 'Nonparametric mixture of regression models', *Journal of the American Statistical Association* **108**(503), 929–941.

Hunter, D. R. & Lange, K. (2004), 'A tutorial on mm algorithms', *The American Statistician* **58**(1), 30–37.

Hunter, D. R. & Young, D. S. (2012), 'Semiparametric mixtures of regressions', *Journal of Nonparametric Statistics* **24**(1), 19–38.

Hyvarinen, A., Karhunen, J. & Oja, E. (2002), 'Independent component analysis', *Studies in Informatics and Control* **11**(2), 205–207.

Kaufman, L. & Rousseeuw, P. J. (2009), *Finding groups in data: an introduction to cluster analysis*, Vol. 344, John Wiley & Sons.

Lee, T.-W., Lewicki, M. S. & Sejnowski, T. J. (1999), 'Unsupervised classification with non-gaussian mixture models using ICA', *Advances in neural information processing systems* pp. 508–514.

Lee, T.-W., Lewicki, M. S. & Sejnowski, T. J. (2000), 'ICA mixture models for unsupervised classification of non-gaussian classes and automatic context switching in blind signal separation', *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **22**(10), 1078–1089.

Levine, M., Hunter, D. R. & Chauveau, D. (2011), 'Maximum smoothed likelihood for multivariate mixtures', *Biometrika* **98**(2), 403–416.

Li, J., Ray, S. & Lindsay, B. G. (2007), 'A nonparametric statistical approach to clustering via mode identification.', *Journal of Machine Learning Research* **8**(8), 1687–1723.

MacQueen, J. et al. (1967), Some methods for classification and analysis of multivariate observations, *in* 'Proceedings of the fifth Berkeley symposium on mathematical statistics and probability', Vol. 1, California, USA, pp. 281–297.

Mallapragada, P. K., Jin, R. & Jain, A. (2010), Non-parametric mixture models for clustering, *in* 'Structural, Syntactic, and Statistical Pattern Recognition', Springer, pp. 334–343.

McLachlan, G. & Peel, D. (2004), *Finite mixture models*, John Wiley & Sons.

Meng, X.-L. & Rubin, D. B. (1993), 'Maximum likelihood estimation via the ECM algorithm: A general framework', *Biometrika* **80**(2), 267–278.

Palmer, J. A., Makeig, S., Kreutz-Delgado, K. & Rao, B. D. (2008), Newton method for the ICA mixture model., *in* 'ICASSP', Citeseer, pp. 1805–1808.

R Core Team (2015), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: *http://www.R-project.org/*

Raykar, V. C., Yang, C., Duraiswami, R. & Gumerov, N. (2005), 'Fast computation of sums of gaussians in high dimensions'.

Sejnowski, T. J. (1999), 'ICA mixture models for image processing'.

Shah, C. A., Arora, M. K. & Varshney, P. K. (2004), 'Unsupervised classification of hyperspectral data: an ICA mixture model based approach', *International Journal of Remote Sensing* **25**(2), 481–487.

Titterington, D. M., Smith, A. F. M. & Makov, U. E. (1985), *Statistical analysis of finite mixture distributions*, John Wiley & Sons.

Vandekerkhove, P. (2013), 'Estimation of a semiparametric mixture of regressions model', *Journal of Nonparametric Statistics* **25**(1), 181–208.

Vichi, M. (2008), 'Fitting semiparametric clustering models to dissimilarity data', *Advances in Data Analysis and Classification* **2**(2), 121–161.

Viele, K. & Tong, B. (2002), 'Modeling with mixtures of linear regressions', *Statistics and Computing* **12**(4), 315–330.

Wolfe, J. H. (1963), Object cluster analysis of social areas, PhD thesis, University of California.

Zhang, W., Fan, J. & Sun, Y. (2009), 'A semiparametric model for cluster data', *Annals of statistics* **37**(5A), 2377.

Zhu, X. & Hunter, D. R. (2015), Theoretical grouding for estimation in conditional independence multivariate finite mixture models. Under review.